

Some Simulation Results for Emphatic Temporal-Difference Learning Algorithms*

Huizhen Yu[†]

Abstract

This is a companion note to our recent study of the weak convergence properties of constrained emphatic temporal-difference learning (ETD) algorithms from a theoretic perspective. It supplements the latter analysis with simulation results and illustrates the behavior of some of the ETD algorithms using three example problems.

Contents

1	About this Note	2
2	Two Test Problems	2
3	Simulation Results for the Constant-stepsize Case	6
3.1	Problem I	6
3.2	Problem II	12
4	Simulation Results for the Diminishing-stepsize Case	18
4.1	Problem I	18
4.2	Problem II	21
5	Mountain Car	23
5.1	Experimental Setup	23
5.2	Simulation Results	26

*This note uses colors in most of the figures to distinguish between different algorithms. Therefore it is better to display the contents on a computer screen than to print them in black and white.

[†]RLAI Lab, Department of Computing Science, University of Alberta, Canada (janey.hzyu@gmail.com). This research was supported by a grant from Alberta Innovates – Technology Futures.

1 About this Note

This is a companion note to our recent study of the weak convergence properties of constrained ETD algorithms from a theoretical perspective [3]. Our purpose here is to supplement that theoretical analysis with simulation results, and to illustrate the behavior of some of the ETD algorithms using examples.

We will consider three test problems: two small grid world-like problems and then the larger Mountain Car problem. As to the algorithms, we will focus on the two variant algorithms in [3] (given by Eqs. (3.3) and (3.4) respectively in Section 3.2 of [3]), as well as their perturbed versions for the constant-stepsize case (given by Eq. (3.7) in Section 3.3 of [3]). These algorithms are constrained ETD algorithms that have biases (cf. the discussion in Section 3.2 of [3]), but they are more robust than the unbiased algorithms in practice, as we will also explain later in Section 2 of this note. We will refer to the two variant algorithms as Variant I and Variant II below.

In what follows, we first describe the two small test problems and illustrate the behavior of the trace iterates (Section 2). We then show simulation results of the constrained ETD algorithms just mentioned, for the case of constant stepsize (Section 3) and for the case of diminishing stepsize (Section 4). We use these results in particular to demonstrate some of the convergence properties proved in [3], and to show that the algorithms are well-behaved despite the high variance issue in off-policy learning. Finally, we show simulation results on the Mountain Car example for a chosen target policy (Section 5). This is to demonstrate that ETD can be applied beyond small test problems and is a useful method for off-policy learning.

Before proceeding, we would like to clarify that we do not intend this note to be a stand-alone paper. We will thus use the notation given in [3] without redefining it here. We will also include very few references – only those needed in order to clarify some experimental setup or results. (Please see the paper [3] for important prior works on TD and ETD learning.)

We would also like to mention that we use colors in most of the figures to distinguish between the iterates produced by different algorithms. Therefore it is better to view the contents of this note on a computer screen than to have them printed out in black and white.

2 Two Test Problems

We now describe two test problems used in our experiments. For these two problems, it is simpler to describe the system dynamics directly in terms of the state transition probabilities, without dealing with actions explicitly. So this is what we are going to do below. (Readers who wish to make the action space explicit may interpret each state transition in our description below as being caused by a distinct action that results in that particular transition with certainty.)

Problem I: Problem I has 6 states. Let P_π and P_{π^o} be the state transition probability matrices under the target policy π and behavior policy π^o , respectively. These transition matrices are given by

$$P_\pi = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0.9 & 0 & 0.1 & 0 & 0 & 0 \\ 0 & 0.9 & 0 & 0.1 & 0 & 0 \\ 0 & 0 & 0.2 & 0.3 & 0.5 & 0 \\ 0 & 0 & 0 & 0.1 & 0 & 0.9 \\ 0.9 & 0 & 0 & 0 & 0.1 & 0 \end{pmatrix}, \quad P_{\pi^o} = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0.5 & 0 & 0.5 & 0 & 0 & 0 \\ 0 & 0.5 & 0 & 0.5 & 0 & 0 \\ 0 & 0 & 0.4 & 0.2 & 0.4 & 0 \\ 0 & 0 & 0 & 0.5 & 0 & 0.5 \\ 0.5 & 0 & 0 & 0 & 0.5 & 0 \end{pmatrix}.$$

Their associated transition graphs have the same topology, which is drawn in Figure 1 (left). The transition from state 6 to 1 has reward 1; all the other transitions have reward zero.

The rest of the parameters are defined as follows. The discount factors $\gamma(s)$ are state-dependent: $\gamma(1) = 0.7$ and $\gamma(s) = 1$ for $s > 1$. The interest weights $i(s)$ and the λ -parameters are also state-

dependent: $i(s) = 1$ for $s \in \{2, 4, 6\}$ and $i(s) = 0$ otherwise, $\lambda(s) = 0$ for $s \in \{2, 4, 6\}$ and $\lambda(s) = 1$ otherwise. Aggregating states into 3 groups, $\{1, 4\}$, $\{2, 3\}$, and $\{5, 6\}$, we assign 3 binary features to each state to indicate its membership.

We remark that with the above choices of i and λ , the approximate value function $\phi(s)^\top \theta^*$ of ETD equals exactly the true value function $v_\pi(s)$ at $s \in \{2, 4, 6\}$, the three states of interest. This serves as an example to show how one can define state-dependent i and λ jointly so that accurate estimates of $v_\pi(s)$ for desired states can be obtained, in spite of capacity limitation in the function approximation architecture.




Figure 1: The transition graphs of two test problems.

Problem II: Problem II has 21 states, whose interconnections are depicted in Figure 1 (right). One state is located at the centre, and the rest of the states split evenly into four groups, indicated by the four loops in Figure 1. The topology of the transition graph is the same for the target and behavior policies. We have drawn the transition graph only for the northeast group in Figure 1 (right); the states in each of the other three groups are arranged in the same manner and have the same transition structure. Given this symmetry, to specify the transition probability matrices P_π and P_{π^o} , it suffices to specify the submatrices of P_π and P_{π^o} for the central state and one of the groups. If we label the central state as state 1 and the states in the northeast group counterclockwise as states 2-6, the submatrices of P_π , P_{π^o} for these states are given by

$$\begin{aligned} \text{target policy:} & \begin{pmatrix} 0 & 0.25 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0.2 & 0 & 0.8 & 0 & 0 \\ 0 & 0 & 0.2 & 0 & 0.8 & 0 \\ 0 & 0 & 0 & 0.2 & 0 & 0.8 \\ 0.8 & 0 & 0 & 0 & 0.2 & 0 \end{pmatrix}, \\ \text{behavior policy:} & \begin{pmatrix} 0 & 0.25 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0.5 & 0 & 0.5 & 0 & 0 \\ 0 & 0 & 0.5 & 0 & 0.5 & 0 \\ 0 & 0 & 0 & 0.5 & 0 & 0.5 \\ 0.5 & 0 & 0 & 0 & 0.5 & 0 \end{pmatrix}. \end{aligned}$$

Intuitively speaking, from the central state, the system enters one group of states by moving diagonally in one of the four directions with equal probability. After spending some time in that group, eventually returns to the central state and the process repeats. The behavior policy on average spends more time wandering inside each group than the target policy, while the target policy tends to traverse counterclockwise through the group more quickly.

All the rewards are zero except for the middle state in each group – for the northeast group, this is the shaded state in Figure 1 (right). For the two northern groups, their middle state has reward 1, while for the two southern groups, that reward is -1 .

The discount factor is $\gamma = 0.9$ for all states. The interest weights and λ -parameters are set to be $\lambda(s) = 0$, $i(s) = 1$ for all states. As to features, we aggregate states into 5 groups, the 4 groups mentioned earlier and the central state forming its own group, and we let each state have 5 binary features indicating its membership.

Behavior of traces: We use the next three figures to illustrate the behavior of traces. (Readers who are interested only in the behavior of the θ -iterates of the ETD algorithms may skip this part and go to the subsequent sections directly.) In general, by identifying certain cycle patterns in the transition graphs, one can infer whether the trace iterates $\{(e_t, F_t)\}$ will be unbounded over time almost surely [2, Section 3.1]. Figure 2 shows a few examples of such cycles in the two test problems just described.




Figure 2: Some cycle patterns in the transition graphs of the test problems.

The left graph in Figure 2 is a cycle of two states, $\{4, 4\}$, in the transition graph of Problem I. The edge of the graph is labeled with the importance sampling weight $0.3/0.2$ for the self-transition $4 \rightarrow 4$. (For the two test problems the importance sampling weights are simply given by the ratios between the entries of P_π and P_{π° .) If we multiply together the importance sampling weight and the discount factor $\gamma(s)$ along this cycle, we get $\frac{0.3}{0.2} \cdot \gamma(4) = \frac{0.3}{0.2} > 1$, while the interest weight $i(4) > 0$ for the only state involved in this cycle. Then from [2, Prop. 3.1] (cf. Footnote 3 therein) we can deduce that in Problem I, the follow-on traces $\{F_t\}$ (which is updated according to $F_t = \gamma_t \rho_{t-1} F_{t-1} + i(S_t)$ in this test problem) will be almost surely unbounded.

Similarly, the right graph in Figure 2 is a cycle of states in the transition graph of Problem II. It consists of the central state and the northeast group of states. The importance sampling weights for each transition are labeled on the edges of the cycle. Traversing through the cycle once from any starting state, and multiplying together the importance sampling weights and the discount factors of each edge and its destination state, we get $\left(\frac{0.8}{0.5}\right)^4 \cdot \gamma^6 = \left(\frac{0.8}{0.5}\right)^4 \cdot 0.9^6 > 1$, while at least one of the states in the cycle has a positive interest weight (since all the states are of interest in this problem). Then we can deduce as in the previous case that $\{F_t\}$ will be unbounded almost surely. Hence, the eligibility traces e_t will also be unbounded in this case (because with $\lambda = 0$ in this problem, we have $e_t = F_t \phi(S_t)$).

As another example, suppose in Problem I we let $\lambda = 1$ for all states instead. The middle graph in Figure 2 exhibits a cycle of states in the transition graph in this case. If we multiply together the importance sampling weights and the $\gamma(s), \lambda(s)$ values on each edge and its destination state in this cycle, we get $\left(\frac{0.9}{0.5}\right)^3 \cdot 0.7 > 1$, while $i(s)\phi(s)$ is nonnegative for all states in the cycle and nonzero for at least one (e.g., state 4 or 6). Then it can be deduced by using [2, Prop. 3.1] as before that the eligibility traces $\{e_t\}$ (generated in this case by $e_t = \lambda_t \gamma_t \rho_{t-1} e_{t-1} + i(S_t)\phi(S_t)$) will be almost surely unbounded.

We plotted in the upper left graphs of Figure 3 and Figure 4 the values of the max-norm $\|(e_t, F_t)\|$ over 8×10^5 iterations for the two test problems, respectively (the x -axis indicates the iteration t). One can see the recurring spikes in these graphs and the exceptionally large values of some of these spikes. This is consistent with the unboundedness of $\{(e_t, F_t)\}$ in the two test problems just discussed.

The unboundedness of $\{(e_t, F_t)\}$ tells us that the invariant probability measure ζ of the Markov chain $\{Z_t\} = \{(S_t, A_t, e_t, F_t)\}$ has an unbounded support. Despite this unboundedness, $\{(e_t, F_t)\}$ is




Figure 3: Statistics of traces for Problem I. See the text for details.




Figure 4: Statistics of traces for Problem II. See the text for details.

bounded in probability (see the discussion in [3, Appendix A]), and under the invariant distribution ζ , $\mathbb{E}_\zeta[\|(e_0, F_0)\|] < \infty$ (see [3, Theorem 2.3]). The latter relation implies that under the invariant distribution, the probability of $\|(e_0, F_0)\| > x$ decreases as $o(1/x)$ for large x . Since the empirical distribution of $\{Z_t\}$ converges to ζ almost surely, during a run of many iterations, we expect to see the fraction of traces with $\|(e_t, F_t)\| > x$ drop in a similar way as x increases.

The simulation results shown in the right part of Figures 3-4 agree with the preceding discussion. Plotted in those two graphs are fractions of traces with $\|(e_t, F_t)\| > x$ during 8×10^5 iterations of the ETD algorithm (the vertical axis indicates the fraction, and the horizontal axis indicates x). For instance, the fraction of traces with $\|(e_t, F_t)\| > 50$ is less than (abound) 0.02 for Problem I (Problem II). It can be seen that despite the recurring spikes in $\|(e_t, F_t)\|$ during the entire run, the fraction of traces with large magnitude x drops sharply with the increase in x .

Finally, let us discuss yet another behavior of the traces, in connection with the biased constrained ETD algorithms that we will focus on in the rest of this note. Although only a small fraction of traces have exceptionally large magnitude, they can occur in consecutive iterations. We plotted two histograms in the lower left part of Figures 3-4 to illustrate this type of behavior. These histograms concern the excursions of the trajectory (e_t, F_t) , $t \geq 0$, outside of the box $\{x \in \mathbb{R}^{n+1} \mid \|x\| \leq 50\}$. The x-axis of the histograms indicates how long is such an excursion (i.e., the number of iterations it contains), and the y-axis indicates how many excursions of length x occurred during the 8×10^5

iterations of the experimental run. We plotted the histograms for length $x > 10$.

This type of behavior suggests that it is better to apply the biased constrained ETD algorithms instead of the unbiased ones (constrained or unconstrained) in practice. This is because when traces with large magnitude occur in consecutive iterations, they can result in large changes in the θ -iterates in a short period of time, if little constraint is put on the size of the change $\theta_{t+1} - \theta_t$ at each iteration. The unbiased ETD algorithms tend to be fragile in practice for this reason, despite their superior asymptotic convergence properties. The biased algorithms take measures to prevent such abrupt changes in the θ -iterates: for example, Variant I truncates the traces, and Variant II truncates the increments in the θ -iterates. Because the fractions of traces with large magnitude are small, these truncations, with proper choices of threshold parameters, make only a small change in the mean update of ETD (cf. the discussion in [3, Section 3.2]). So the biased algorithms can gain much robustness in performance by paying only a small price of bias.

3 Simulation Results for the Constant-stepsize Case

In this section we show simulation results of the biased constrained ETD algorithms with a constant stepsize, for the two test problems described in the previous section. Besides the two biased algorithms, Variant I and Variant II, we will also show results for the perturbed versions of these two variants. Our focus will be on the behavior of multiple consecutive θ -iterates and the behavior of a trajectory of θ -iterates or their averaged iterates, under various stepsizes.

In the experiments reported below, the radius parameter r_B for constraining θ is set to be $r_B = 100$ (well above the threshold required by [3, Lemma 2.1], which is calculated to be $r_B > 7.04$ for Problem I and $r_B > 5.20$ for Problem II). The function ψ_K in the variant algorithms (cf. Eq. (3.2) in [3]) is taken to be the componentwise truncation, $\psi_K(x) = \min\{K, \max\{-K, x\}\}$, for $K = 50$. The perturbed versions of the two algorithms use the same r_B and ψ_K , and the perturbation variables $\Delta_{\theta,t}^\alpha$ (which are of the same size as θ) are i.i.d. normal random variables with zero mean and covariance matrix $(\frac{\alpha}{2})^2 I$.

We will also show results of a modified version of ELSTD, which like Variant I also uses ψ_K to truncate the traces e_t in its matrix/vector iterates. The limiting θ produced by this modified ELSTD is indeed the point that Variant I would converge to in the case of diminishing stepsize. Thus by running this version of ELSTD we can get an estimate of the bias in Variant I. To be concise, in what follows, we will often refer to this modified ELSTD algorithm simply as ELSTD.

To visualize the behavior of the algorithms, instead of plotting the iterates θ_t themselves, we will calculate and plot the normalized distances between θ_t and the desired ETD solution θ^* . Here by the normalized distance we mean $|\theta_t - \theta^*|/|\theta^*|$, normalized by $|\theta^*|$, which is nonzero for both test problems. Correspondingly, we will refer frequently to δ -neighborhoods of θ^* where δ are multiples of $|\theta^*|$, such as the $0.1|\theta^*|$ -neighborhood of θ^* or the $x|\theta^*|$ -neighborhood of θ^* for some $x > 0$.

3.1 Problem I

The experiments below compare the behavior of the various algorithms in Problem I, for four different stepsizes: $\alpha = 0.01, 0.002, 0.001, 0.0005$. First, we did 4 independent runs of both Variant I and Variant II. Each run lasted for 6×10^5 iterations, during which the same state trajectory is used by both algorithms for all the four stepsizes. We did the same experiments for the perturbed versions of the two variants. To illustrate the steady state behavior, we used only the last 4×10^5 iterations of each run to obtain the statistics of multiple consecutive iterates shown in Figures 6-9 below.

Before proceeding to explain these figures, let us first show an example trajectory from a single run (more trajectories of iterates will be shown later). Plotted in Figure 5 are the normalized distances to θ^* of the iterates θ_t^α produced by Variants I and II (top row) and by their perturbed versions (bottom row) in the last 4×10^5 iterations of one run, for the smallest stepsize $\alpha = 0.0005$.

in our experiments. The dashed lines correspond to the averaged iterates $\bar{\theta}_t^\alpha$ (where the averaging also starts with the later portion of the run and neglects its initial portion). It can be seen that compared to the original iterates θ_t^α , the averaged iterates $\bar{\theta}_t^\alpha$ are much less volatile and, in the case of the unperturbed variant algorithms, approach a smaller neighborhood of θ^* .

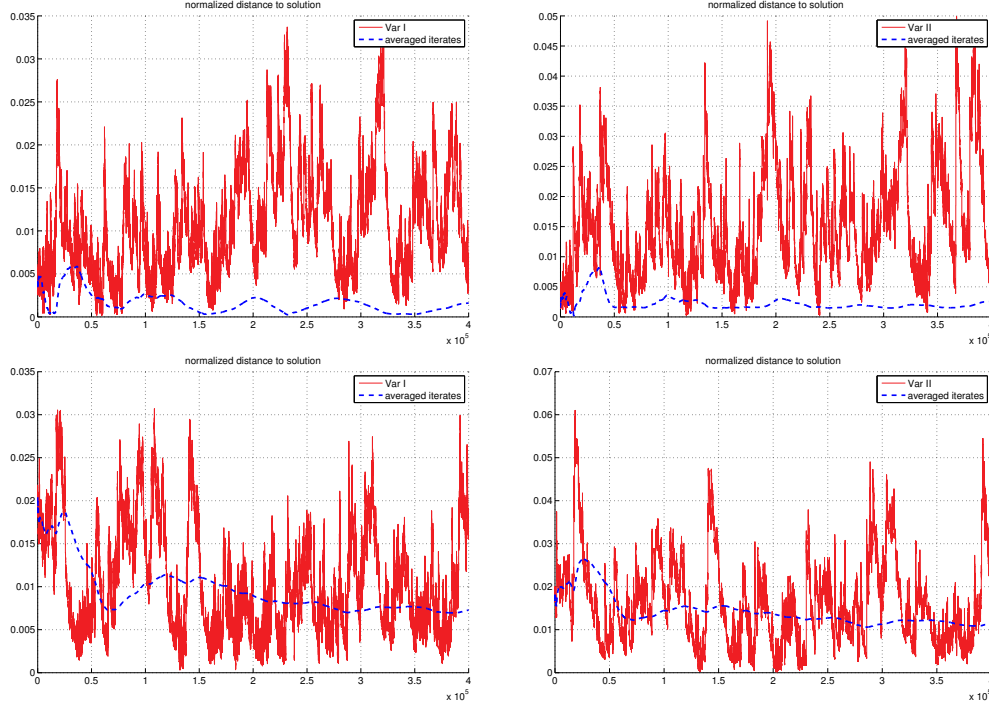


Figure 5: Variants I and II without (top) and with (bottom) perturbation ($\alpha = 0.0005$).

Another note is that since the algorithms are biased, even if we had used smaller stepsizes, the iterates would not be able to approach an arbitrarily small neighborhood of θ^* . To get an estimate of the degree of bias, we ran the modified ELSTD for 8 independent runs of 8×10^5 iterations each. Averaged over the 8 runs, the mean normalized distance (to θ^*) of the ELSTD final solution was 0.0035 with standard deviation 0.0017. Consistently, we see from Figure 5 that most iterates of Variants I and II are still outside the $0.005|\theta^*|$ -neighborhood of θ^* , although the averaged iterates $\bar{\theta}_t^\alpha$ seem to approach a smaller neighborhood. Recall also that Variant II need not converge at all (cf. [3, Section 3.2]). In our experiments we observed it to behave similarly to Variant I and have a comparable bias (albeit slightly larger than that of Variant I for the two test problems).

We now proceed to explain the details of Figures 6-9, which demonstrate the behavior of multiple consecutive θ -iterates for the four stepsizes:

Figures 6-7: In both figures, the x -axis represents the $x|\theta^*|$ -neighborhood of θ^* , and the y -component of a point (x, y) represents the fraction of times that a certain number of consecutive iterates θ_t^α fail to lie entirely inside the $x|\theta^*|$ -neighborhood of θ^* . Specifically, for Figure 6 we consider every segment of 100 consecutive iterates, $(\theta_t^\alpha, \dots, \theta_{t+99}^\alpha)$, $t = 0, 1, \dots$, during each run. Plotted in Figure 6 are the fractions of times (during a single run) that such a segment fails to lie entirely inside the $x|\theta^*|$ -neighborhood of θ^* . We then consider segments of $\lfloor \frac{1}{\alpha} \rfloor$ consecutive iterates, $(\theta_t^\alpha, \dots, \theta_{t+\lfloor 1/\alpha \rfloor - 1}^\alpha)$, $t = 0, 1, \dots$. Plotted in Figure 7 are the fractions of times (during a single run) that a segment of length $\lfloor \frac{1}{\alpha} \rfloor$ fails to lie entirely inside the $x|\theta^*|$ -neighborhood of θ^* . (Note that the smaller the stepsize α , the longer the segments used to calculate the fractions of times shown in the figure.) In both figures, for each color and each algorithm, the solid line corresponds to the results from one of the four runs,

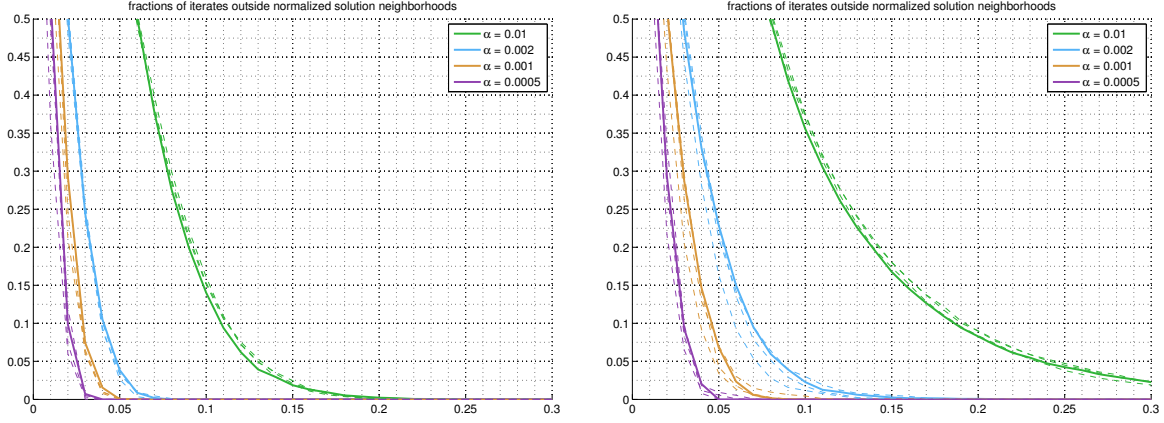


Figure 6: Variant I (left) and Variant II (right) without perturbation. The x -axis represents the $x|\theta^*|$ -neighborhood of θ^* . The y -component of a point (x, y) represents the fraction of times (in a single run) that a segment of 100 consecutive iterates fails to lie entirely inside the $x|\theta^*|$ -neighborhood of θ^* . (See the text for more details.)




Figure 7: Variant I (left) and Variant II (right) without perturbation. The x -axis represents the $x|\theta^*|$ -neighborhood of θ^* . The y -component of a point (x, y) represents the fraction of times (in a single run) that a segment of $\lfloor \frac{1}{\alpha} \rfloor$ consecutive iterates fails to lie entirely inside the $x|\theta^*|$ -neighborhood of θ^* . (See the text for more details.)

while the three dashed lines correspond to the results from the other three runs. It can be seen that the smaller the stepsize, the smaller the neighborhood of θ^* inside which a trajectory of iterates spends most of its time. The behavior of multiple consecutive iterates shown in these figures can be compared with the assertions in Theorem 3.4(ii) and Theorem 3.6(i) of [3].

Figures 8-9: We repeated the same experiments for the perturbed versions of Variants I and II. The results are shown in Figures 8-9, and they show similar behavior of the multiple consecutive iterates generated by these perturbed algorithms. (As in the previous case, in the figures, for each color, the solid lines correspond to the results from one of the four runs, and the dashed lines the other three runs.) These simulation results can be compared with the assertions in Theorem 3.8 of [3].

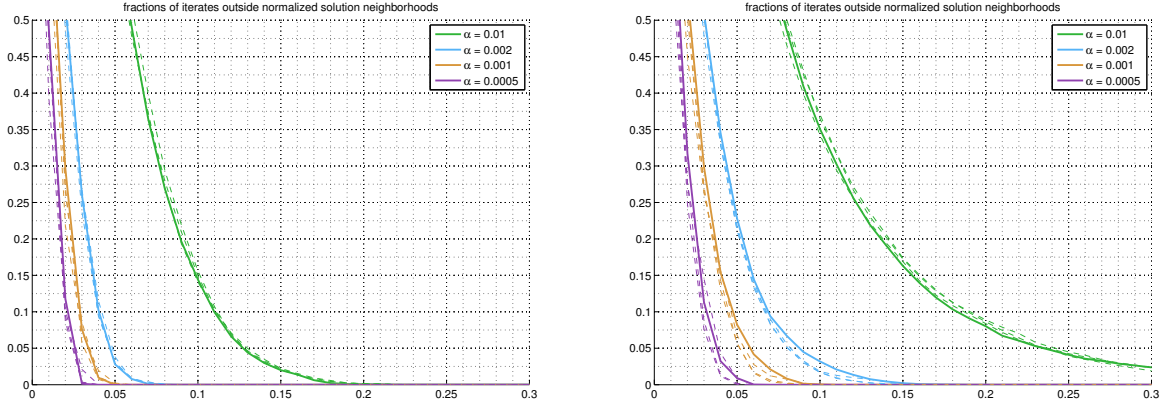


Figure 8: Variant I (left) and Variant II (right) with perturbation. The x -axis represents the $x|\theta^*|$ -neighborhood of θ^* . The y -component of a point (x, y) represents the fraction of times (in a single run) that a segment of 100 consecutive iterates fails to lie entirely inside the $x|\theta^*|$ -neighborhood of θ^* . See the text for more details.

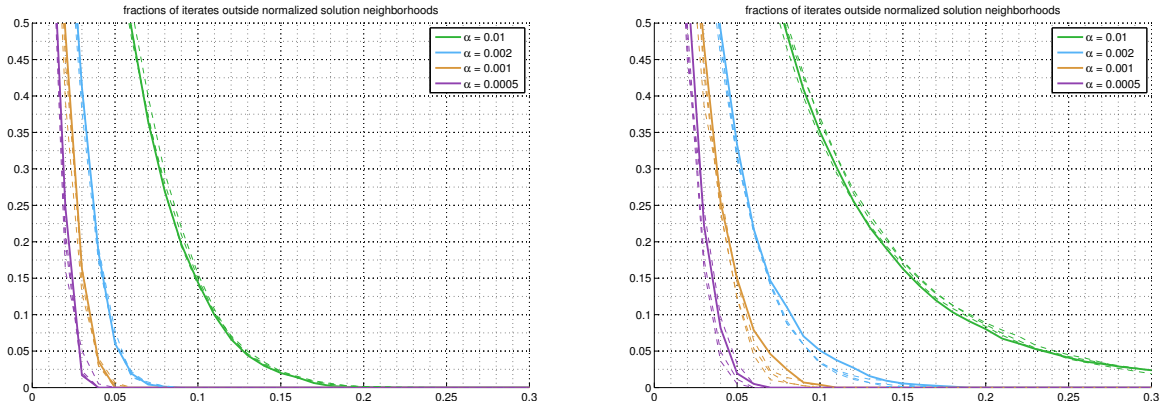


Figure 9: Variant I (left) and Variant II (right) with perturbation. The x -axis represents the $x|\theta^*|$ -neighborhood of θ^* . The y -component of a point (x, y) represents the fraction of times (in a single run) that a segment of $\lfloor \frac{1}{\alpha} \rfloor$ consecutive iterates fails to lie entirely inside the $x|\theta^*|$ -neighborhood of θ^* . (See the text for more details.)

In the rest of this subsection we show more trajectories of iterates from individual runs. The results are plotted in Figures 10-13, and the details of the experiments and our observations from them are as follows.

Figures 10-11: In these plots we show the normalized distances (to θ^*) of a trajectory of averaged iterates $\bar{\theta}_t^\alpha$ and original iterates θ_t^α , for each algorithm and each stepsize, using the data from one of the experimental runs that produced the previous four figures. Comparing the top rows with the bottom rows in Figures 10-11, the averaged iterates $\bar{\theta}_t^\alpha$ are better than θ_t^α in terms of both the volatility of the iterates and the closeness to the desired solution, especially when the stepsize is relatively large. Comparing the right columns with the left ones in Figures 10-11, it can be seen that Variant II has a larger variance than Variant I (although we have also observed the opposite in other problems not reported in this note). Comparing Figure 10 with Figure 11, it can be seen that for the same stepsize, the perturbed algorithms settled inside a larger neighborhood of θ^* than the unperturbed algorithms did. This suggests that the better asymptotic properties of the perturbed



Figure 10: Variant I (left) and Variant II (right) without perturbation. Top: averaged iterates $\bar{\theta}_t^\alpha$; bottom: iterates θ_t^α . Data are from a single run.

algorithms can be compromised by the noises brought by the perturbation (cf. Remark 3.2 at the end of Section 3.3 of [3]), and the unperturbed algorithms may be adequate for practical purposes (cf. Remark 4.1 at the end of Section 4.3 in [3]).

Figure 12: In this experiment we compare the transient behavior of the variant algorithms for the four stepsizes using a single run of 10^5 iterations. All the algorithms start from the same initial condition, and no portion of the run is discarded. ELSTD is also included for comparison: the linear equations formed by ELSTD are solved every 500 iterations to produce the ELSTD curve shown in the figure. It can be seen that ELSTD converges rapidly. With a large stepsize $\alpha = 0.01$, Variants I and II also make quick initial progress, before they start to oscillate in a relatively large neighborhood of θ^* .

Figure 13: This experiment serves as an example to show that the convergence behavior of the ETD algorithms are not affected when the matrix C associated with ETD is negative semidefinite instead of negative definite (cf. [3, Section 5.1]). In this experiment we let $i(s) = 1$ for only two states, $s \in \{2, 6\}$ (i.e., removing state 4 from the original list of states of interest) and we set $i(s) = 0$ for the other states. Correspondingly, we set $\lambda(s) = 0$ for $s \in \{2, 6\}$ and $\lambda(s) = 1$ otherwise. Then the 3×3 matrix C has rank 2 and becomes negative semidefinite. We ran the unperturbed Variant I and Variant II, initialized at zero, and we also ran ELSTD. The algorithms ran as in the previous cases and none of them had any issues (cf. the explanations given in Section 5.1 of [3]). The iterates have higher variances in this case, so in order to obtain iterates that can approach the $0.1|\theta^*|$ -neighborhood of θ^* , we used a larger threshold $K = 200$ in the function ψ_K , as well as

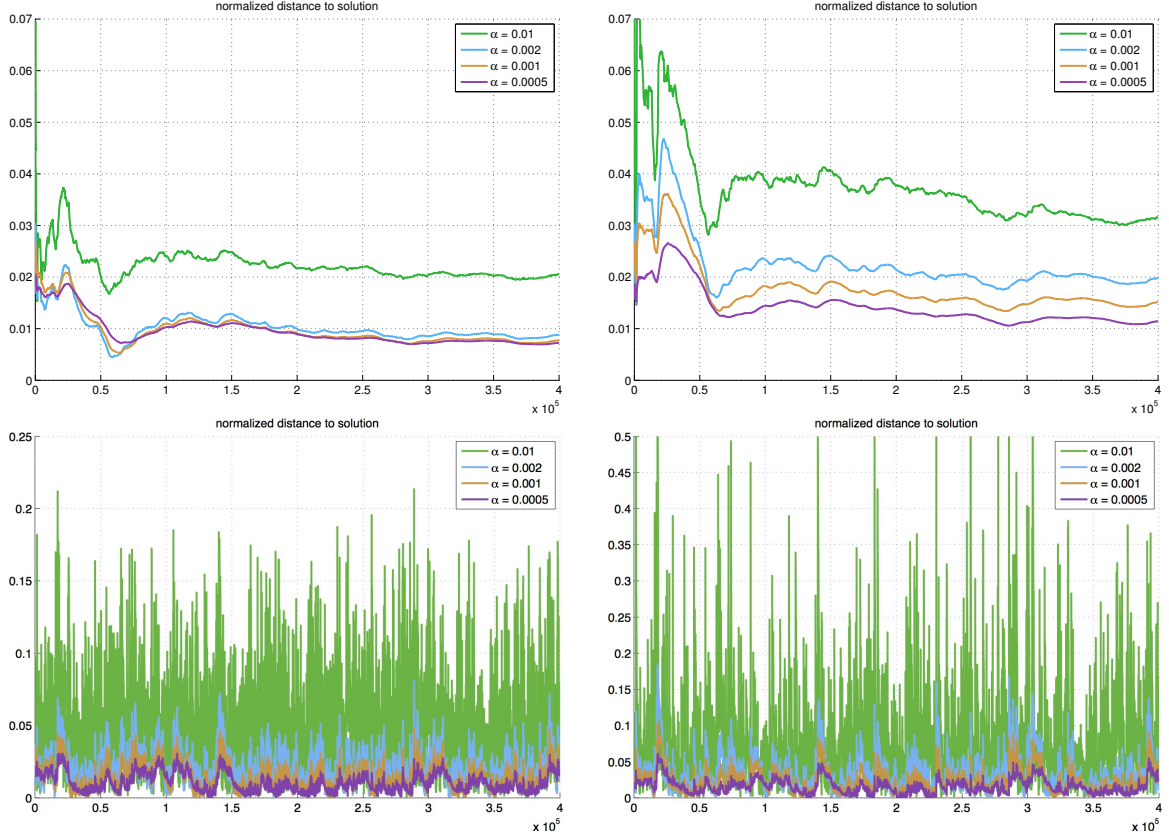


Figure 11: Variant I (left) and Variant II (right) with perturbation. Top: averaged iterates $\bar{\theta}_t^\alpha$; bottom: iterates θ_t^α . Data are from a single run.

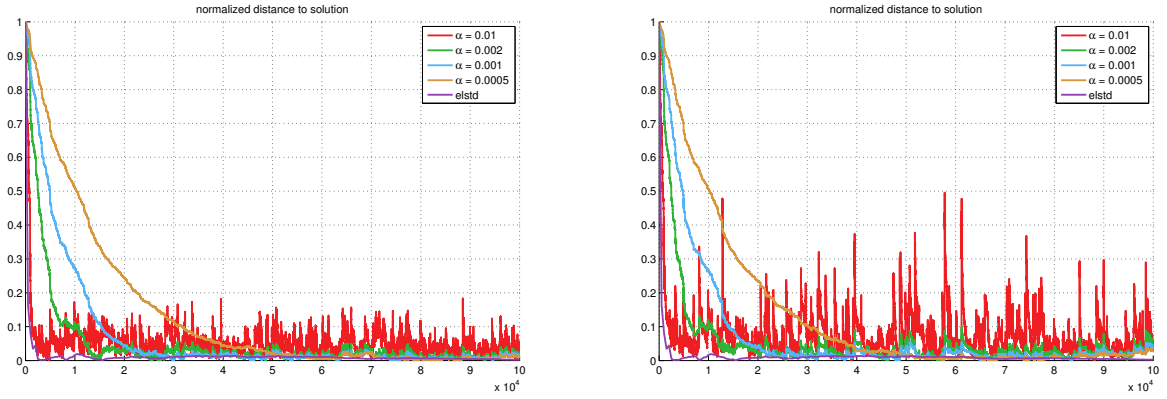


Figure 12: Variant I (left) and Variant II (right) without perturbation. Data are from a single run; ELSTD is also included for comparison.

smaller stepsizes in this experiment. (The higher variances in this case have nothing to do with the negative semidefiniteness of C . Instead it is a consequence of the following fact: here ETD is solving a generalized multistep Bellman equation for the two states $\{2, 6\}$ of interest, and these two states are far apart from each other in the directed transition graph. So compared with the original setting

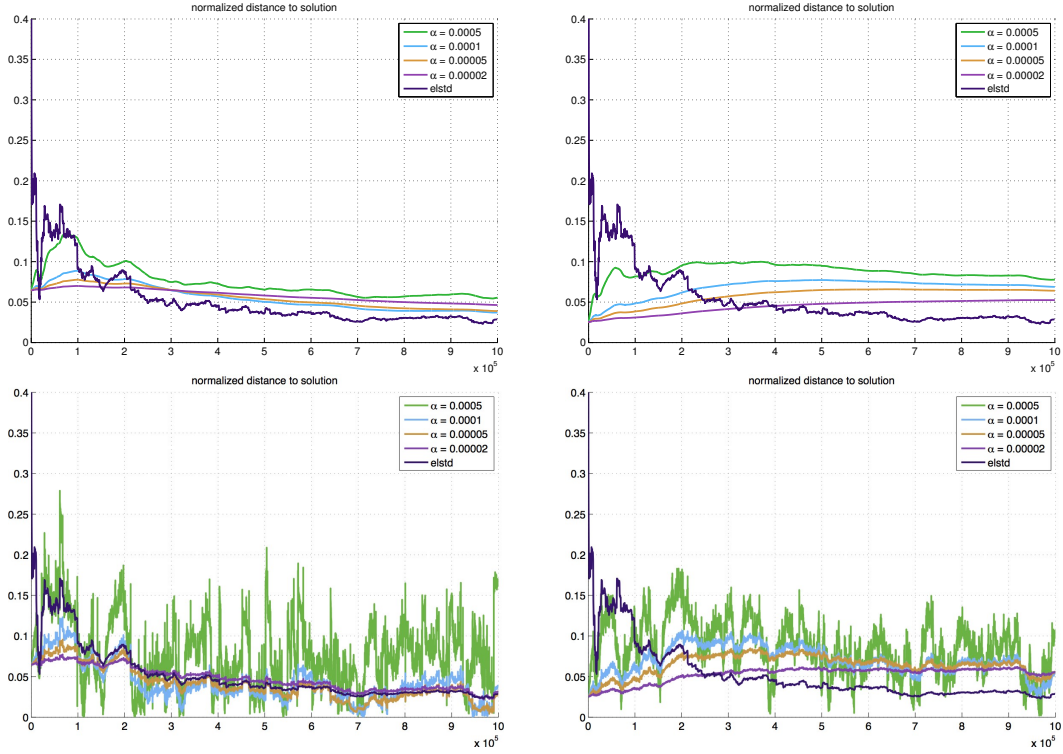


Figure 13: Variant I (left) and Variant II (right) without perturbation. Top: averaged iterates $\bar{\theta}_t^\alpha$; bottom: iterates θ_t^α . Data are from a single run; ELSTD is included in all the cases for comparison.

of Problem I, in this case it takes on average more steps to reach any of the states of interest again after visiting either one of them.)

Plotted in Figure 13 are the normalized distances to θ^* of the averaged iterates $\bar{\theta}_t^\alpha$ and of the original iterates θ_t^α generated in the later portion of a single run, for four choices of stepsizes. Specifically, to reduce transient effects and focus on the steady state behavior, we first ran all the algorithms for 3×10^5 iterations with the stepsize 0.0005, and we then continued the run for another 10^6 iterations with the four different stepsizes indicated in the figure. The averaged iterates shown in the top row of the figure are generated from those later 10^6 iterations of the run. It can be seen that overall the behavior of the iterates is similar to what we observed in the previous experiments.

3.2 Problem II

We repeated for Problem II the same experiments we did for Problem I. All the algorithms are tested for five stepsizes: $\alpha = 0.0005, 0.0002, 0.0001, 0.00005, 0.00002$. First, we did 4 independent runs of Variants I and II and their perturbed versions. Each run has 11×10^5 iterations, and the last 8×10^5 iterations are used to obtain the statistics of multiple consecutive iterates shown in Figures 15-18, in order to show the steady state behavior of the algorithms. More details are as follows.

Figure 14: This figure shows an example trajectory from a single run. Plotted in the figure are the normalized distances to θ^* of the iterates θ_t^α generated by the four algorithms for the smallest stepsize $\alpha = 0.00002$. The dashed lines correspond to the averaged iterates $\bar{\theta}_t^\alpha$, which, like in the case of Problem I, can be seen to behave better than the original iterates θ_t^α . We ran ELSTD to get an estimate of the degree of bias of Variant I. Averaged over 8 independent runs of 8×10^5 iterations each, the mean normalized distance of the ELSTD final solution to θ^* was 0.043 with




Figure 14: Variants I and II without (top) and with (bottom) perturbation ($\alpha = 0.00002$).

standard deviation 0.003. So based on Figure 14, it seems that the iterates generated by Variant I with the stepsize $\alpha = 0.00002$ are not far from the smallest neighborhood that Variant I can reach.

Figures 15-18: These figures show the behavior of multiple consecutive iterates for the four algorithms with different stepsizes. For Variants I and II, plotted in Figure 15 are the fractions of times (during a single run) that a segment of length 100, $(\theta_t^\alpha, \dots, \theta_{t+99}^\alpha)$, fails to lie entirely inside the $x|\theta^*|$ -neighborhood, and plotted in Figure 16 are the fractions of times (during a single run) that a segment of length $\lfloor \frac{1}{\alpha} \rfloor$, $(\theta_t^\alpha, \dots, \theta_{t+\lfloor 1/\alpha \rfloor - 1}^\alpha)$, fails to lie entirely inside the $x|\theta^*|$ -neighborhood. For the perturbed version of Variant I and Variant II, the same plots are shown in Figure 17 and Figure 18, respectively. In all these figures, for each color and each algorithm, the solid lines correspond to the results from one of the four runs, and the dashed lines the other three runs. The behavior exhibited here is similar to what we observed in the case of Problem I: as the stepsize becomes smaller, the iterates spread out less and the trajectory spends more time inside a smaller neighborhood of θ^* . This can be compared with the assertions in Theorem 3.4(ii) and Theorem 3.6(i) of [3] for Variants I and II, and with the assertions in Theorem 3.8 of [3] for the perturbed versions of these two variants.

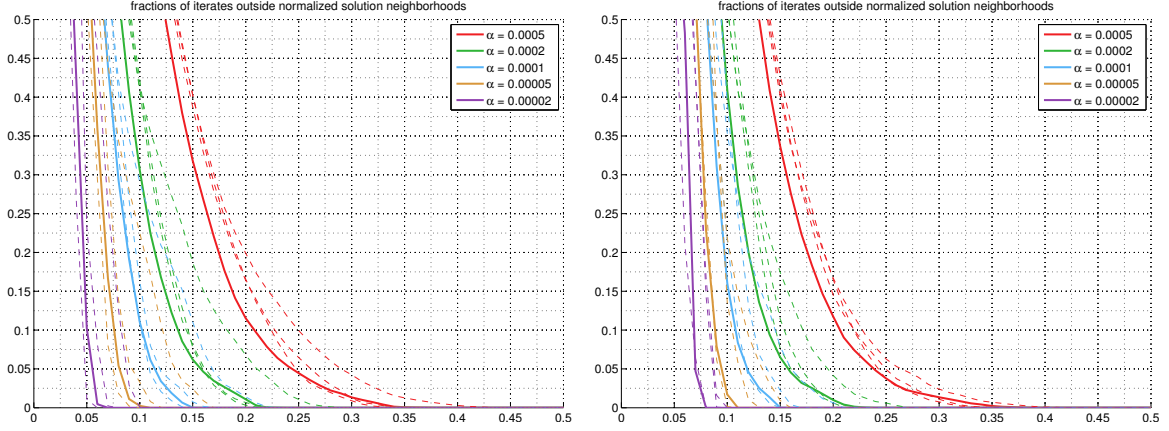


Figure 15: Variant I (left) and Variant II (right) without perturbation. The x -axis represents the $x|\theta^*|$ -neighborhood of θ^* . The y -component of a point (x, y) represents the fraction of times (in a single run) that a segment of 100 consecutive iterates fails to lie entirely inside the $x|\theta^*|$ -neighborhood of θ^* .

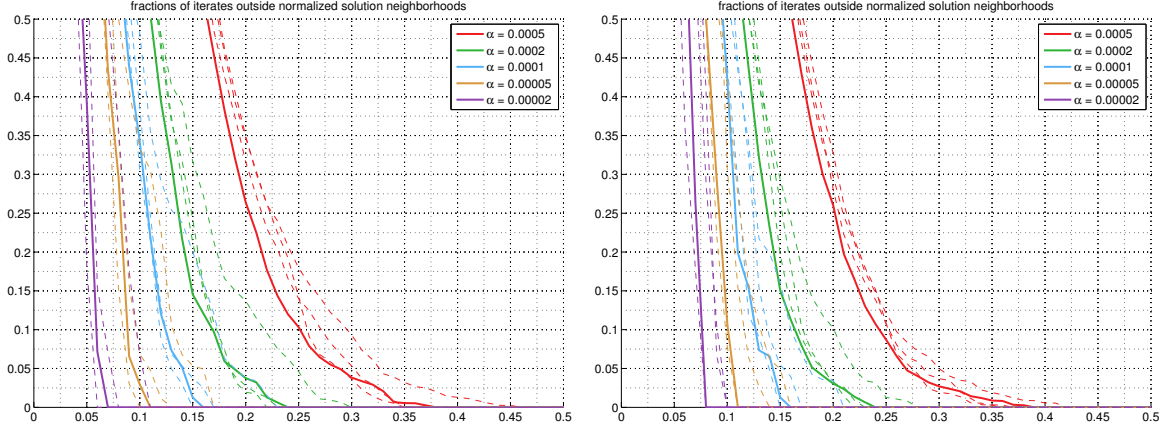


Figure 16: Variant I (left) and Variant II (right) without perturbation. The x -axis represents the $x|\theta^*|$ -neighborhood of θ^* . The y -component of a point (x, y) represents the fraction of times (in a single run) that a segment of $\lfloor \frac{1}{\alpha} \rfloor$ consecutive iterates fails to lie entirely inside the $x|\theta^*|$ -neighborhood of θ^* .




Figure 17: Variant I (left) and Variant II (right) with perturbation. The x -axis represents the $x|\theta^*|$ -neighborhood of θ^* . The y -component of a point (x, y) represents the fraction of times (in a single run) that a segment of 100 consecutive iterates fails to lie entirely inside the $x|\theta^*|$ -neighborhood of θ^* .




Figure 18: Variant I (left) and Variant II (right) with perturbation. The x -axis represents the $x|\theta^*|$ -neighborhood of θ^* . The y -component of a point (x, y) represents the fraction of times (in a single run) that a segment of $\lfloor \frac{1}{\alpha} \rfloor$ consecutive iterates fails to lie entirely inside the $x|\theta^*|$ -neighborhood of θ^* .

In the rest of this subsection we show more trajectories of iterates from individual runs. The details of the experiments are as follows.

Figures 19-20: In these two figures we plotted the normalized distances to θ^* of a trajectory of averaged iterates $\bar{\theta}_t^\alpha$ and original iterates θ_t^α , for each algorithm and each stepsize, using the data from one of the runs of the algorithms that produced the previous four figures. Our observations from these results are the same as those from Figures 10-11 in the case of Problem I: (i) the averaged iterates $\bar{\theta}_t^\alpha$ perform better than θ_t^α in that they vary less and can approach a smaller neighborhood of θ^* ; (ii) the unperturbed algorithms do not seem to have any disadvantages compared with the perturbed algorithms for the same stepsize.

Figure 21: This experiment compares the transient behavior of the variant algorithms using a single run of 6×10^5 iterations. All the algorithms start from the same initial condition, no portion of the run is discarded, and ELSTD is also included for comparison. The linear equations formed by ELSTD are solved every 500 iterations to produce the ELSTD curve shown in the figure. It can be seen that ELSTD converges rapidly. The variant algorithms can make fast initial progress too with the largest stepsize $\alpha = 0.0005$, although because of the big stepsize, they quickly start to oscillate in a relatively large neighborhood of θ^* .

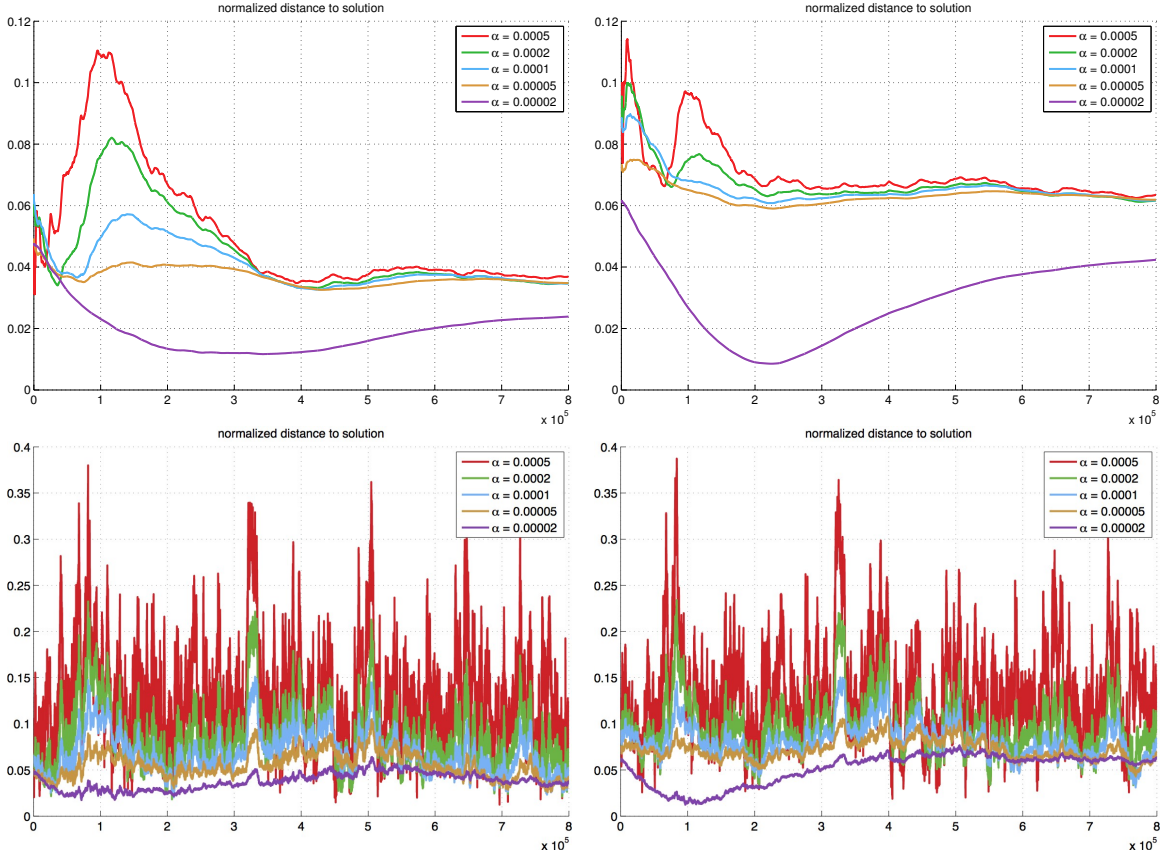


Figure 19: Variant I (left) and Variant II (right) without perturbation. Top: averaged iterates $\bar{\theta}_t^\alpha$; bottom: iterates θ_t^α . Data are from a single run.

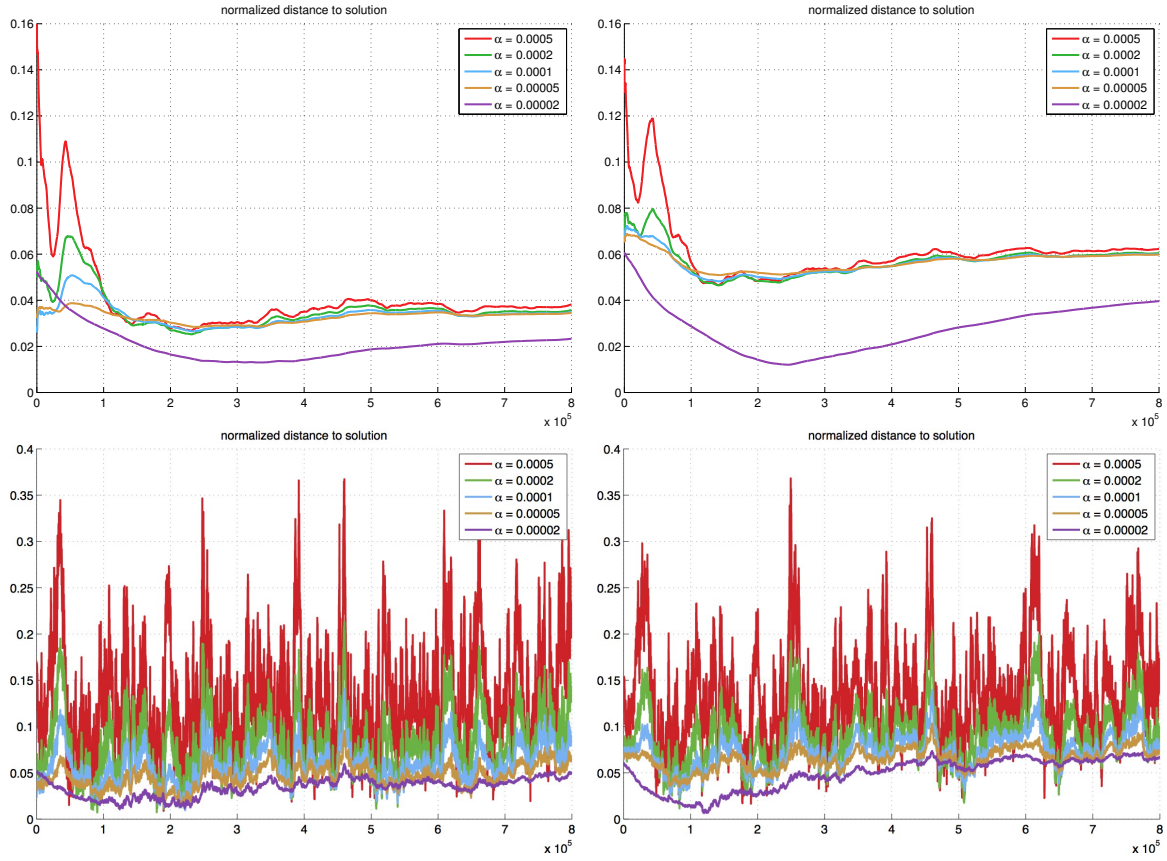


Figure 20: Variant I (left) and Variant II (right) with perturbation. Top: averaged iterates $\bar{\theta}_t^\alpha$; bottom: iterates θ_t^α . Data are from a single run.




Figure 21: Variant I (left) and Variant II (right) without perturbation. Data are from a single run; ELSTD is also included for comparison.

4 Simulation Results for the Diminishing-stepsize Case

In this section we illustrate the behavior of Variant I and Variant II with diminishing stepsize for the two test problems. As in the previous case, we set the radius parameter $r_B = 100$ and use the componentwise truncation function ψ_K with $K = 50$ in the two variant algorithms. For visualizing the behavior of the θ -iterates as well as the behavior of multiple consecutive iterates, we will plot their normalized distances to the desired ETD solution θ^* as before. Given the close connection between the constant-stepsize case and the diminishing-stepsize case, and given also what we already observed in the former case, the results from the present part of the experiments, to be reported below, turn out to be as expected.

4.1 Problem I

In the first experiment, we used five stepsize sequences that decrease at different rates β ,

$$\alpha_t = \frac{1}{200 + (0.1t)^\beta} \quad \text{for } \beta \in \{0.3, 0.5, 0.7, 0.9, 1\}.$$

We ran the two algorithms with these five stepsize rules simultaneously for 6×10^5 iterations, using a common state trajectory. The results are plotted in Figure 22. ELSTD (modified as in Section 3) is also included for comparison: the linear equations formed by ELSTD are solved every 500 iterations to produce the ELSTD curve in the figure.

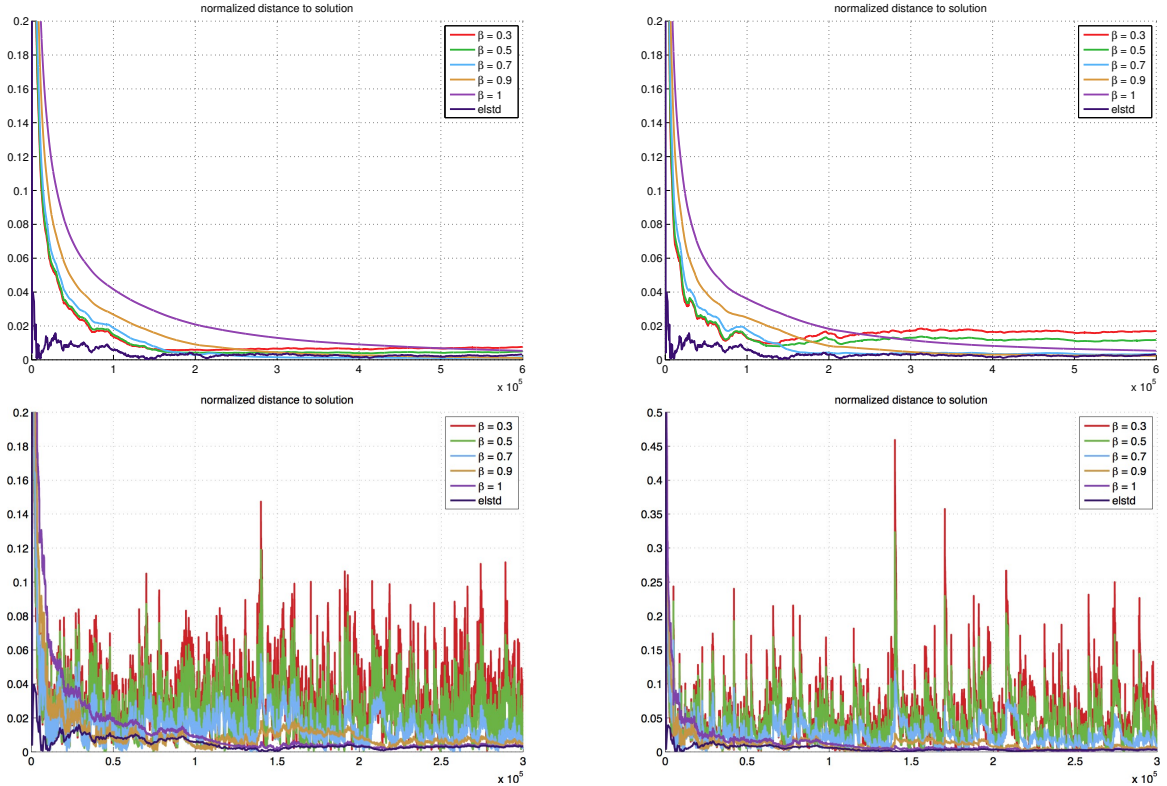


Figure 22: Variant I (left) and Variant II (right). Top: averaged iterates $\bar{\theta}_t$ for the entire run; bottom: iterates θ_t for the first half of the run. ELSTD is also included for comparison. See the text for details.

The top row of Figure 22 shows the normalized distances of the averaged iterates $\bar{\theta}_t$ for the entire run, and the bottom row shows the normalized distances of the iterates θ_t for only the first half of the run, in order to have a close-up view of the transient behavior. Comparing the top row with the bottom row, the advantages of the averaged iterates for large stepsizes, especially $\beta = 0.3, 0.5$, can be seen. (The use of averaging for $\beta < 1$ is known as Polyak-averaging). We can also see that the iterates θ_t for $\beta = 0.3$ or 0.5 did not settle in a small neighborhood of θ^* like the iterates generated with smaller stepsizes. This can be explained as follows: Even after $t = 6 \times 10^5$ iterations, we have $\alpha_t \approx 0.004$ for $\beta = 0.3$ and $\alpha_t \approx 0.002$ for $\beta = 0.5$, so we can expect the iterates for $\beta = 0.3$ or 0.5 to behave at best like the iterates with constant stepsize 0.002 (cf. the bottom row of Figure 10).

The next experiment was designed in accordance with the observation of the relation between the diminishing-stepsize case and the constant-stepsize case just mentioned. We did 10 independent runs of 10^6 iterations each, for the two variant algorithms using two stepsize rules:

$$\alpha_t = \frac{1}{200 + (5t)^\beta} \quad \text{for } \beta = 0.7, \quad \text{and} \quad \alpha_t = \frac{1}{200 + (200t)^\beta} \quad \text{for } \beta = 0.5.$$

Since we want to test the convergence behavior of the algorithms, in defining the preceding stepsize rules, we have made sure that the stepsize becomes small enough later in the run (at $t = 10^6$, α_t is of the order 10^{-5} in both cases of β). The simulation results are plotted in Figure 23 and Figure 24 for $\beta = 0.7$ and $\beta = 0.5$ respectively. The details are as follows.

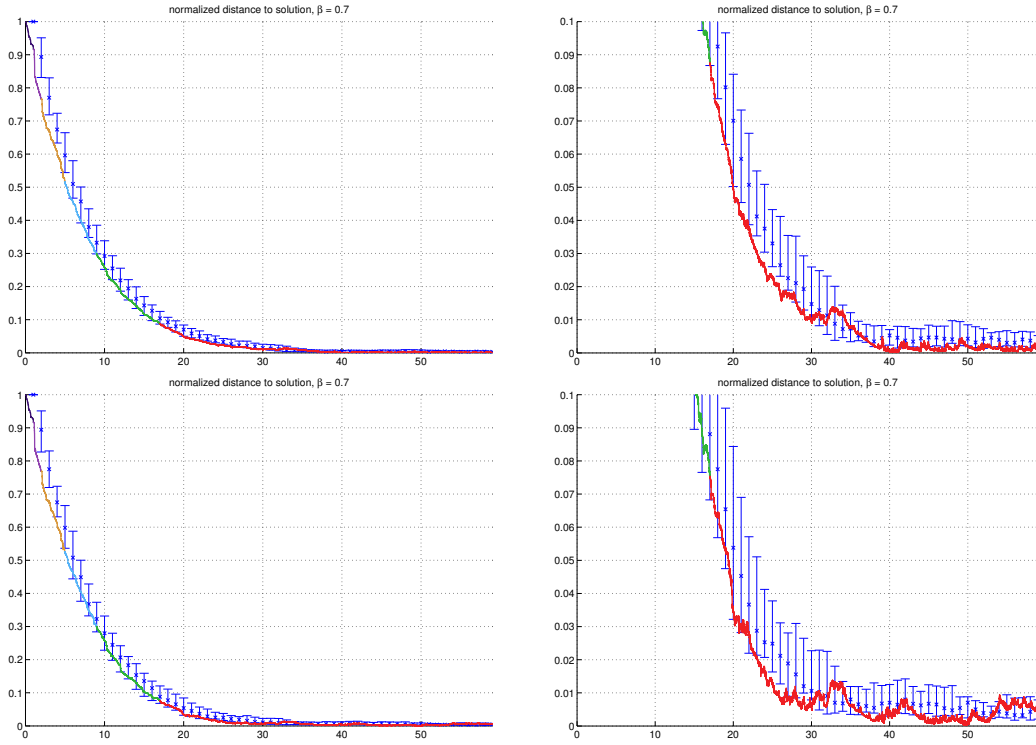


Figure 23: Variant I (top) and Variant II (bottom) with $\beta = 0.7$. The bottom portion of each plot on the left is enlarged and shown on the right. The solid curve corresponds to one run, with the x -component being $x = \sum_{k=0}^t \alpha_k$ and the y -component being the normalized distance of θ_t to θ^* . (See the text for more details.)

In Figures 23-24, plotted in solid lines are the normalized distances (to θ^*) of the iterates from one of the 10 runs. Specifically, each solid curve is made up of points $(\sum_{k=0}^t \alpha_k, |\theta_t - \theta^*|/|\theta^*|)$,

$t \geq 0$, from a single run of an algorithm. In words, the x -axis represents a continuous timeline (cf. [3, Section 3.1]), and the x -component of a solid curve corresponds to the sum of stepsizes up to an iteration, whereas the y -component of the curve corresponds to the normalized distance of that iterate. The whole curve is plotted on the left side of each figure, with a close-up view of its bottom portion shown on the right side.

To give a rough indication of the values of the decreasing stepsizes themselves, we colored segments of the solid curves in different colors according to the range of stepsizes in each segment as follows: $\alpha_t \geq 0.003$ (black), $\alpha_t \in (0.003, 0.002]$ (purple), $\alpha_t \in (0.002, 0.001]$ (brown), $\alpha_t \in (0.001, 0.0005]$ (blue), $\alpha_t \in (0.0005, 0.0002]$ (green), $\alpha_t < 0.0002$ (red).

The blue error bars in Figures 23-24 give statistics about the maximal deviation from θ^* for multiple consecutive iterates from the 10 experimental runs. The horizontal positions of these error bars equal positive integers x , and for each x , the x -th error bar is generated as follows. For each run, the iterates θ_t are grouped into segments such that the x -th segment consists of those θ_t with $\sum_{k=0}^t \alpha_k \in [x-1, x)$. So as x increases, the x -th segment consists of more and more iterates, but measured with respect to the continuous timeline, all the segments are of length 1 approximately. We then calculate for each $x \geq 1$ the maximal normalized distance for the θ -iterates in the x -th segment of each run, $\max_{x\text{-th segment}} |\theta_t - \theta^*|/|\theta^*|$. This gives us 10 numbers, one for each run. We take the median, min and max of these numbers to form the x -th error bar. In particular, the point with an ‘ \times ’ mark inside the bar is the median, and the lower and upper ends of the bar correspond to the minimum and maximum of the 10 numbers, respectively.

The simulation results shown in Figures 23-24 can be compared with the assertion in Theorem 3.3 of [3] for Variants I and II with diminishing stepsizes.




Figure 24: Variant I (top) and Variant II (bottom) with $\beta = 0.5$. The bottom portion of each plot on the left is enlarged and shown on the right. The solid curve corresponds to one run, with the x -component being $x = \sum_{k=0}^t \alpha_k$ and the y -component being the normalized distance of θ_t to θ^* . (See the text for more details.)

4.2 Problem II

Similar to the previous subsection, in the first experiment for Problem II, we used five stepsize sequences that decrease at different rates β :

$$\alpha_t = \frac{1}{2000 + (0.1t)^\beta} \quad \text{for } \beta \in \{0.3, 0.5, 0.7, 0.9, 1\}.$$

We ran the two algorithms with these five stepsize rules simultaneously for 8×10^5 iterations, using a common state trajectory. The results are plotted in Figure 25. ELSTD (modified as in Section 3) is also included for comparison: the linear equations formed by ELSTD are solved every 500 iterations to produce the ELSTD curve in the figure. The top row of Figure 22 shows the normalized distances of the averaged iterates $\bar{\theta}_t$ for the entire run, and the bottom row shows the normalized distances of the iterates θ_t only for the first half of the run, in order to have a close-up view of the transient behavior. Once more, for $\beta < 1$, the advantages of averaging can be seen.

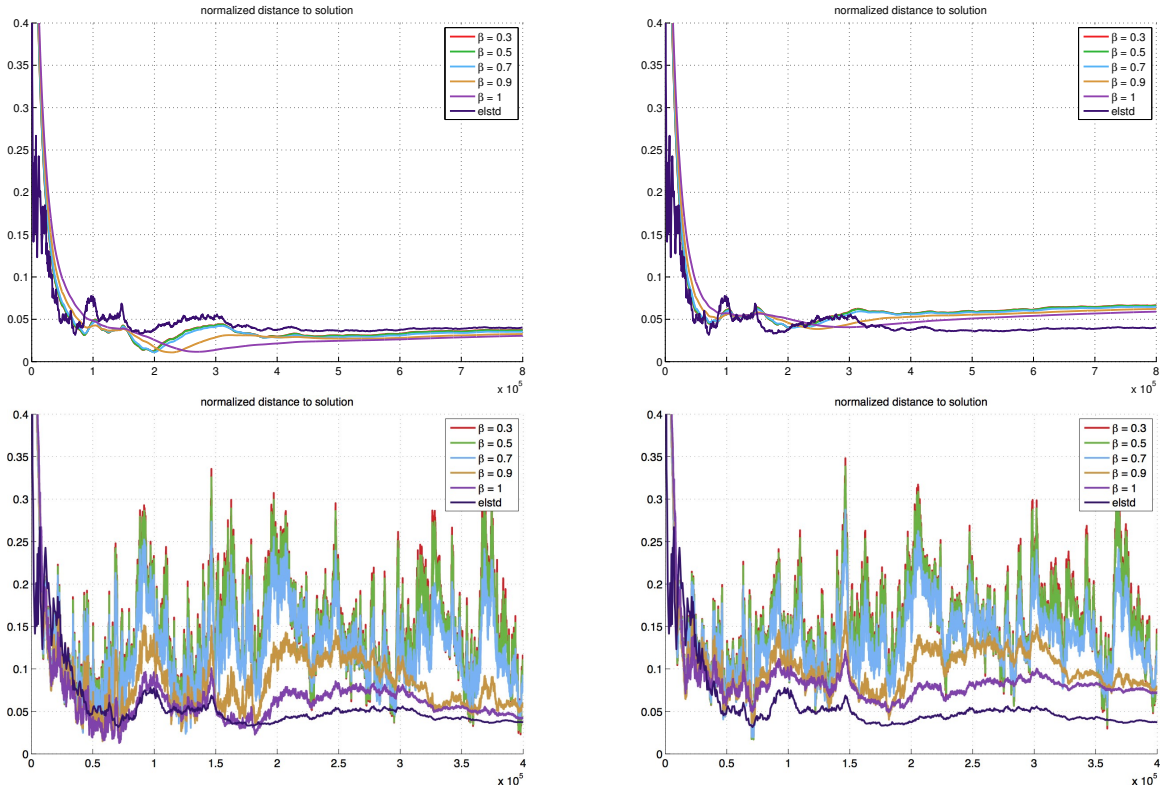


Figure 25: Variant I (left) and Variant II (right). Top: averaged iterates $\bar{\theta}_t$ for the entire run; bottom: iterates θ_t for the first half of the run. ELSTD is also included for comparison.

It can be seen that for the three largest stepsize rules ($\beta \in \{0.3, 0.5, 0.7\}$), the iterates behaved similarly to each other and did not settle in a small neighborhood of θ^* like the iterates generated with smaller stepsizes. This can again be understood by relating the situation here to the constant-stepsize case: With $\beta \in \{0.3, 0.5, 0.7\}$, the stepsizes decrease rather slowly. Even after $t = 8 \times 10^5$ iterations, α_t is between 0.0004 and 0.0005 for $\beta = 0.3, 0.5$ and about 0.0002 for $\beta = 0.7$. So we can expect the iterates to behave at best like the iterates with constant stepsize 0.0002 (cf. the bottom row of Figure 19).

In the next experiment, we did 10 independent runs of 1.6×10^6 iterations each, for the two variant algorithms using two stepsize rules:

$$\alpha_t = \frac{1}{2000 + (10t)^\beta} \quad \text{for } \beta = 0.7, \quad \text{and} \quad \alpha_t = \frac{1}{2000 + (4000t)^\beta} \quad \text{for } \beta = 0.5.$$

As before we chose these stepsize rules to ensure that the stepsize becomes small enough later in the run (at $t = 1.6 \times 10^6$, α_t is about 10^{-5} in both cases of β). The simulation results are plotted in Figure 26 and Figure 27 for $\beta = 0.7$ and $\beta = 0.5$ respectively. The graphical objects in these figures have the same meanings as those in Figures 23-24 for Problem I, so we will only describe these objects briefly here.

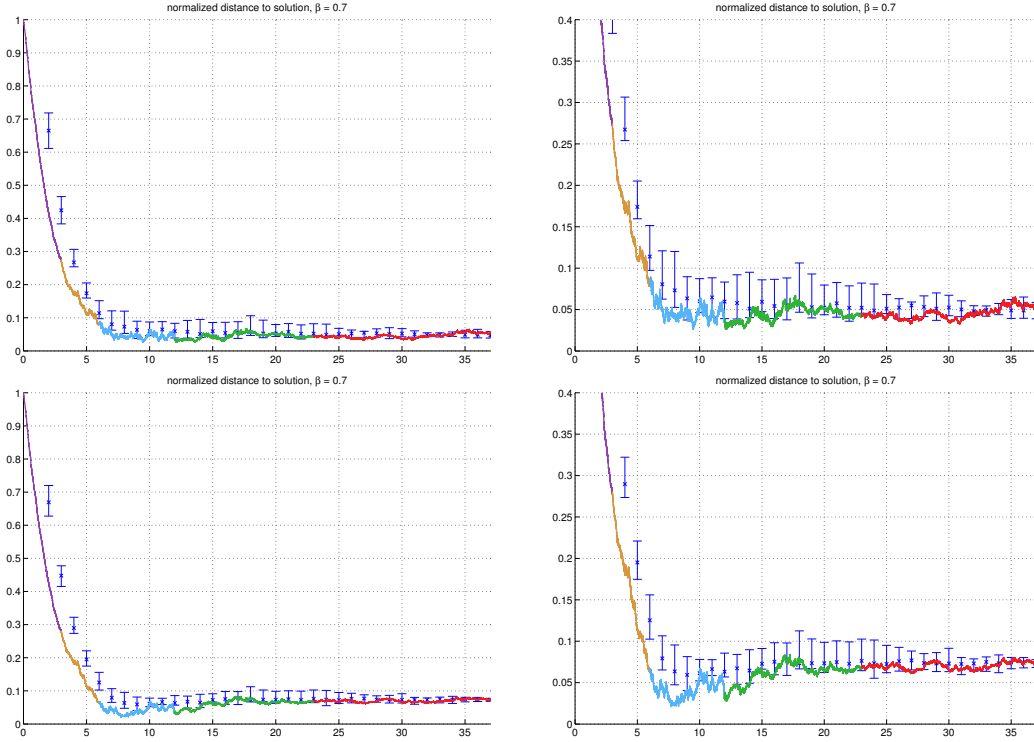


Figure 26: Variant I (top) and Variant II (bottom) with $\beta = 0.7$. The bottom portion of each plot on the left is enlarged and shown on the right. The solid curve corresponds to one run, with the x -component being $x = \sum_{k=0}^t \alpha_k$ and the y -component being the normalized distance of θ_t to θ^* . (See the text for more details.)

In Figures 26-27, we plotted in solid lines the normalized distances (to θ^*) of the iterates from one of the 10 runs. The x -axis represents a continuous timeline, and a solid curve is made up of points $(\sum_{k=0}^t \alpha_k, |\theta_t - \theta^*|/|\theta^*|)$ from a single run of an algorithm. The whole curve is plotted on the left side of each figure, with a close-up view of its bottom portion shown on the right side. We colored segments of the curves in different colors according to the range of stepsizes in each segment as follows: $\alpha_t \in (0.0005, 0.0002]$ (purple), $\alpha_t \in (0.0002, 0.0001]$ (brown), $\alpha_t \in (0.0001, 0.00005]$ (blue), $\alpha_t \in (0.00005, 0.00002]$ (green), $\alpha_t < 0.00002$ (red).

The blue error bars in the figures show the range of the maximal deviation from θ^* for multiple consecutive iterates from the 10 experimental runs. They are formed in the same way as described in the earlier experiment for Problem I (see the descriptions for Figures 23-24 in the previous subsection). The point with an ‘ \times ’ mark inside the x -th error bar is the median, and the lower

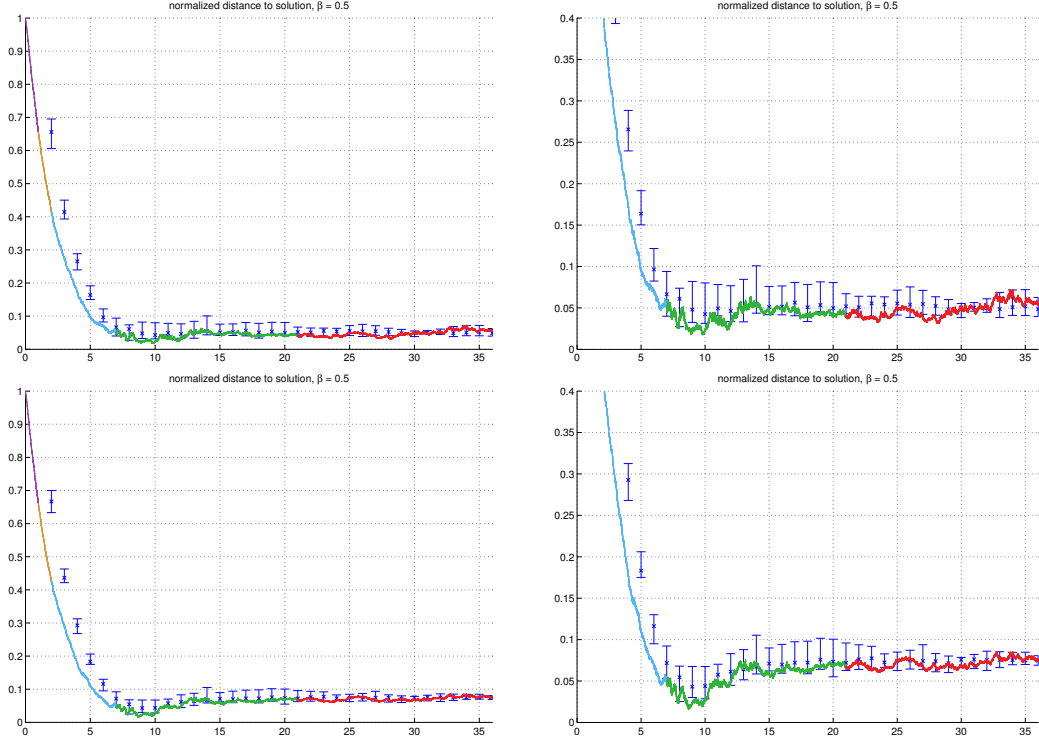


Figure 27: Variant I (top) and Variant II (bottom) with $\beta = 0.5$. The bottom portion of each plot on the left is enlarged and shown on the right. The solid curve corresponds to one run, with the x -component being $x = \sum_{k=0}^t \alpha_k$ and the y -component being the normalized distance of θ_t to θ^* . (See the text for more details.)

and upper ends of the error bar are the minimum and maximum, respectively, of the 10 values of $\max_{x\text{-th segment}} |\theta_t - \theta^*|/|\theta^*|$ obtained from the 10 independent experimental runs.

The simulation results shown in Figures 23-24 can be compared with the assertion in Theorem 3.3 of [3] for the two variant algorithms with diminishing stepsizes.

5 Mountain Car

In this last set of experiments we test constrained ETD on a larger problem constructed from the Mountain Car problem [1]. Mountain Car has continuous state and action spaces. As such it is actually beyond the finite-space model considered in [3], so the convergence theorems we proved therein for constrained ETD do not extend to the Mountain Car problem. Nevertheless, we observed empirically in our experiments that constrained ETD is well-behaved, and in this section we report some of these simulation results for Variant I with a constant stepsize. (Variant II behaves similarly but with a larger variance for this problem.)

5.1 Experimental Setup

We take the dynamics of the Mountain Car problem. The goal is to drive an underpowered car to reach the top of a steep hill. A state consists of the position p and velocity v of the car, whose values are bounded as $p \in [-1.2, 0.5]$ and $v \in [-0.07, 0.07]$. The position 0.5 corresponds to the desired hill top destination, while the position $-\pi/6$ corresponds to the bottom of a valley. At each state

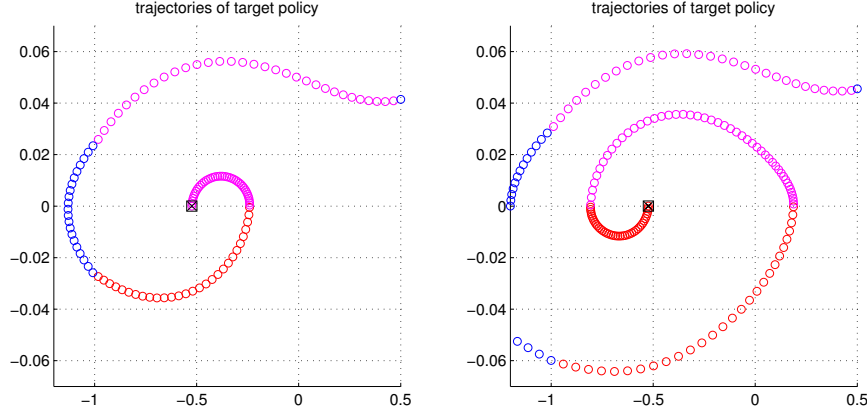


Figure 28: Two state trajectories that can occur (with equal probability 0.5) under the target policy for the initial state $(-\frac{\pi}{6}, 0)$ (indicated by the marked square). The color of a state indicates the action taken at that state: red for **back**, blue for **coast**, and purple for **forward**.

three actions are available: $\{\text{back}, \text{coast}, \text{forward}\}$, designated by $\{-1, 0, 1\}$, respectively. With A_t denoting the action taken at time t and with $\Pi_{[a,b]}(x) = \max\{a, \min\{b, x\}\}$ for an interval $[a, b]$ and scalar x , the dynamics of the car are defined as

$$v_{t+1} = \Pi_{[-0.07, 0.07]}(v_t + 0.001A_t - 0.0025 \cos(3p_t)), \quad p_{t+1} = \Pi_{[-1.2, 0.5]}(p_t + v_{t+1}),$$

except that when $p_{t+1} = -1.2$, the velocity is reset to zero: $v_{t+1} = 0$. Before the destination $p = 0.5$ is reached, the rewards depend only on the action taken and are given by $r(-1) = -1.5$, $r(1) = -1$, and $r(0) = 0$. Once the destination $p = 0.5$ is reached, the car enters a rewardless termination state permanently. We consider undiscounted expected total rewards, so the discount factor $\gamma = 1$.

Target policy: The following policy will be our target policy π throughout the experiments: at a state (p, v) ,

- if $p < -1$, then take action 0 (coast);
- if $p \geq -1$, then take action $\text{sign}(v)$ unless $|v| \leq 10^{-6}$, in which case take action ± 1 (forward or back) with equal probability.

This is a simple policy but behaves reasonably well. Figure 29 (next page) shows the negative value function $-v_\pi$ and its contour map. The values of v_π shown in this figure are estimated by simulating the policy 200 times for each starting state (p, v) in a set of 171×141 points evenly spaced in the position-velocity space $[-1.2, 0.5] \times [-0.07, 0.07]$. In particular, the position (velocity) interval is evenly divided into subintervals of length 0.01 (0.001). Figure 28 above shows the two trajectories that the car can traverse through in the state space if it is initially parked at the bottom of the valley ($p = -\pi/6$) and follows the target policy π .

Behavior policy: We use a fixed sampling scheme to generate states, actions and transitions for ETD learning. This scheme serves the role of the behavior policy π^o and is defined as follows.

- At a state (p, v) , if $p = 0.5$ (the desired destination), then the next state is sampled uniformly from the state space $[-1.2, 0.5] \times [-0.07, 0.07]$.
- For $p \neq 0.5$, two things can happen:
 - (1) With probability 0.9, an action is chosen from the set $\{\text{back}, \text{coast}, \text{forward}\}$ randomly and uniformly, and the next state is determined by the state transition under that action.
 - (2) With probability 0.1, a random state (p', v') is chosen as the next state. In particular,




Figure 29: $-v_\pi$ estimated by simulating π (left: 3D view; right: contour map).

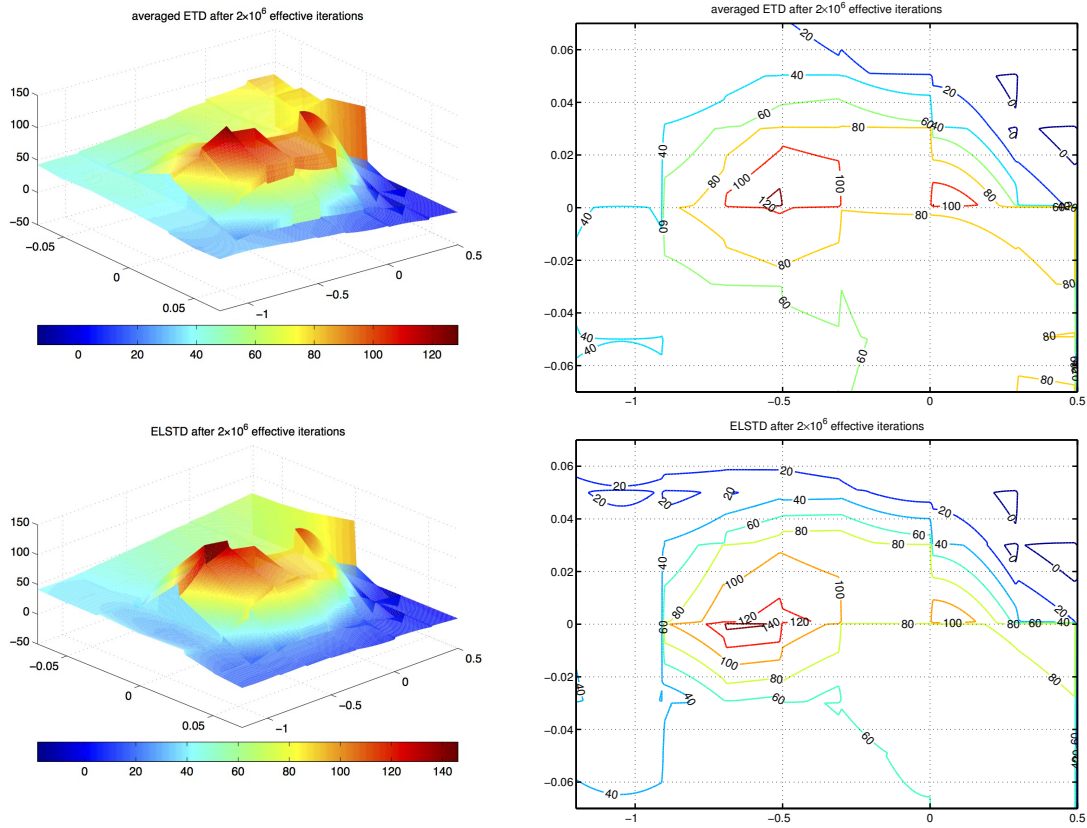


Figure 30: Piecewise linear approximation of $-v_\pi$ (linear in $(\cos(3p), v)$ on each piece) calculated by Variant I (top) and ELSTD (bottom). Left: 3D view; right: contour map.

either the velocity remains the same, $v' = v$, and the position p' is uniformly sampled from the interval $[p, 0.5]$ or $[-1.2, p]$ (each of these two cases happens with probability 0.04), or (p', v') is uniformly sampled from the state space (this happens with probability 0.02).

The above scheme of generating data can be viewed as a valid behavior policy by enlarging the action space to include three more actions that correspond to the three different ways of randomly choosing

(p', v') described in step (2) above. This defines the importance sampling weights $\pi(s, a)/\pi^o(s, a)$ for the constrained ETD algorithms in the experiments.

It is worth mentioning that in the mathematical framework of off-policy learning, we need not restrict the behavior policy to be a physically feasible policy. Indeed, that would limit the use of off-policy learning in the goal-reaching type of problem such as Mountain Car, since in such problems, to find a policy that is able to reach the goal state can be tantamount to solving the problem itself. By defining the behavior policy in a broader way, one can apply off-policy learning methods to solving goal-reaching problems, at least in the context where the system dynamics of the problems can be simulated.

Algorithmic parameters: We will only show results for Variant I with a constant stepsize, as mentioned earlier. The following algorithmic parameters are used throughout the experiments: a constant interest weight 0.5 and a constant $\lambda = 0.5$ for all states; stepsize $\alpha = 0.003$; the radius parameter $r_B = 2 \times 10^4$ for constraining θ ; and the truncation function ψ_K with $K = 50$ as given before. Since our purpose here is only to demonstrate that ETD can be applied beyond synthetic small problems, we did not optimize over these parameters. The stepsize we used is relatively large. As in the previous sections, we find that the use of a larger stepsize can make the algorithm progress faster initially, and together with averaging, it can yield useful approximation results in fewer iterations.

5.2 Simulation Results

First experiment: It can be seen from Figure 29 (previous page) that v_π is discontinuous and can change sharply between certain regions of the state space. In the first experiment, we partition the position (and velocity) interval into 7 (and 6) subintervals to form 42 rectangular regions to cover the space $[-1.2, 0.5] \times [-0.07, 0.07]$. In each region we approximate v_π by an affine function of $(\cos(3p), v)$; the entire approximation is thus piecewise linear in $(\cos(3p), v)$. Specifically, to partition the position interval $[-1.2, 0.5]$ and the velocity interval $[-0.07, 0.07]$, we use the points given in the two vectors below as the mid points:

$$\text{position: } (-0.9 \quad -0.7 \quad -0.5 \quad -0.3 \quad 0 \quad 0.3), \quad \text{velocity: } (-0.05 \quad -0.03 \quad 0 \quad 0.03 \quad 0.05).$$

Each of the 42 regions is the product of two left-closed right-open intervals, except at the boundary of the state space, where the end points of an interval can be included or excluded in order to fill exactly the space $[-1.2, 0.5] \times [-0.07, 0.07]$. For each region, we used these 3 features, 1, $\cos(3p)$, and $15v$, to approximate the value function in that region.

The approximate value function obtained by Variant I after a single run of 2×10^6 effective iterations is shown in the top row of Figure 30 (previous page). Here an *effective iteration* refers to an iteration in which the behavior policy takes an action that could also be taken by the target policy. Plotted is the approximation corresponding to the averaged iterate $\bar{\theta}_t^\alpha$ at the end of the run, where the average is taken over the last 10^6 iterations to avoid transient effects. We also ran ELSTD (modified as before) in the same run for comparison, and the approximate value function it obtained is plotted in the bottom row of Figure 30. It can be seen that both algorithms try to approximate v_π , and overall the contours of their approximations roughly match the contour of v_π in shape (cf. Figure 29).

More experiments: In the subsequent experiments we ran Variant I with features generated by tile-coding [1]. This gives piecewise constant approximations of v_π , where the pieces are defined by the tile-coding schemes we use. So instead of comparing the ETD approximations to v_π , which has jumps and curvy contours as shown in Figure 29, it seems better to compare the ETD approximations to an approximate solution from a discretized model that has a discretization resolution comparable to the resolution of the tile-coding schemes in the experiments. Such a discretized model was




Figure 31: Approximate $-v_\pi$ from a discretized model (left: 3D view; right: contour map).




Figure 32: Piecewise constant approximation of $-v_\pi$ calculated by Variant I (top) and ELSTD (bottom) using tile-coding. Left: 3D view; right: contour map.

constructed and it yields the approximate v_π shown in Figure 31. The ETD approximation results for two tile-coding schemes are shown in Figures 32-34. The details of these figures are as follows.

Figure 31: We first built a discretized finite-state model for the target policy as follows. We divided the position (velocity) interval evenly into subintervals of length 0.1 (0.01), and thus obtained 270 rectangular regions in total to fill the space $[-1.2, 0.5] \times [-0.07, 0.07]$. (As in the first experiment,

except at the boundaries of the state space, each of these regions is the product of two left-closed right-open intervals.) The states in each region is treated as one aggregate state in the discretized model. To define the transition probabilities between the aggregate states for the target policy, we ran the behavior policy (the sampling scheme described earlier) for 10^7 iterations, and used those effective iterations to calculate the transition frequencies between the aggregate states. These frequencies are taken to be the transition probabilities in the discretized model, and the per-stage rewards for the model are defined similarly. The Bellman equation for the discretized model is then solved, and the solution is used to define a piecewise constant approximation of v_π (constant over each aggregate state). Figure 31 plots the 3D view and contour map of the resulting approximation, which may be compared with the estimated v_π shown in Figure 29.

Figure 32: In this experiment a coarse tile-coding scheme is used to generate 78 overlapping rectangular regions in total to cover the state space. Specifically, a first tiling comprises of 36 rectangles, which are obtained by dividing the position interval $[-1.2, 0.5]$ and the velocity interval $[-0.07, 0.07]$ unevenly at the following points:

$$\text{position: } (-0.9 \ -0.6 \ -0.3 \ 0 \ 0.3), \quad \text{velocity: } (-0.05 \ -0.02 \ 0 \ 0.02 \ 0.05).$$

A second tiling is similarly formed by dividing the position and velocity intervals at these points:

$$\text{position: } (-1.0 \ -0.7 \ -0.4 \ -0.1 \ 0.2), \quad \text{velocity: } (-0.06 \ -0.04 \ -0.01 \ 0.01 \ 0.03 \ 0.06).$$

This tiling comprises of 42 rectangles (it corresponds to offsetting the first tiling by $(0.2, 0.01)$ and then covering the exposed sides of the state space with extra rectangles). Correspondingly, we used $36 + 42 = 78$ binary features for each state, to indicate the two rectangles containing that state. (As before, each rectangle is taken to be the product of two left-closed right-open intervals except on the boundaries of the state space.)

We ran Variant I and ELSTD with these features for 10^6 effective iterations. Plotted in the top row of Figure 32 is the approximate value function corresponding to the averaged iterate $\bar{\theta}_t^\alpha$ produced by Variant I at the end of the run. The average here is taken over the last 5×10^5 iterations to reduce transient effects. The approximation obtained by ELSTD is plotted in the bottom row of Figure 32 for comparison. The resolutions of the two tilings used in this experiment are lower than the resolution used to build the discretized model. Nevertheless, comparing Figure 32 with Figure 31, one can recognize the similarities between the ETD/ELSTD approximations here and the approximate v_π from the discretized model shown.

Figures 33-34 (next page): In this experiment we ran Variant I with a finer tile-coding scheme. Similarly to the previous case, we made two tilings of the state space by dividing the position and velocity intervals unevenly, first at these points:

$$\text{position: } (-1.0 \ -0.8 \ -0.6 \ -0.4 \ -0.2 \ 0 \ 0.2), \quad \text{velocity: } (-0.05 \ -0.03 \ -0.01 \ 0 \ 0.01 \ 0.03 \ 0.05),$$

and then at these points:

$$\begin{aligned} \text{position: } & (-1.1 \ -0.9 \ -0.7 \ -0.5 \ -0.3 \ -0.1 \ 0.1 \ 0.3), \\ \text{velocity: } & (-0.06 \ -0.04 \ -0.02 \ 0 \ 0.01 \ 0.02 \ 0.04 \ 0.06). \end{aligned}$$

The first (second) tiling comprises of 64 (81) rectangles. Correspondingly, each state has $64 + 81 = 145$ binary features to indicate the two rectangles that contain the state.

We ran both Variant I and ELSTD with these features. Figure 33 shows the approximations obtained after 2×10^5 effective iterations. As can be seen, Variant I is in the process of building up the approximate value function, while ELSTD generally converges faster. Figure 34 shows the approximations obtained after 10^6 effective iterations, where for Variant I (top row of Figure 34), plotted is the approximate value function corresponding to the averaged iterate $\bar{\theta}_t^\alpha$ at the end of the




Figure 33: Piecewise constant approximation of $-v_\pi$ calculated by Variant I (left) and ELSTD (right) using tile-coding.

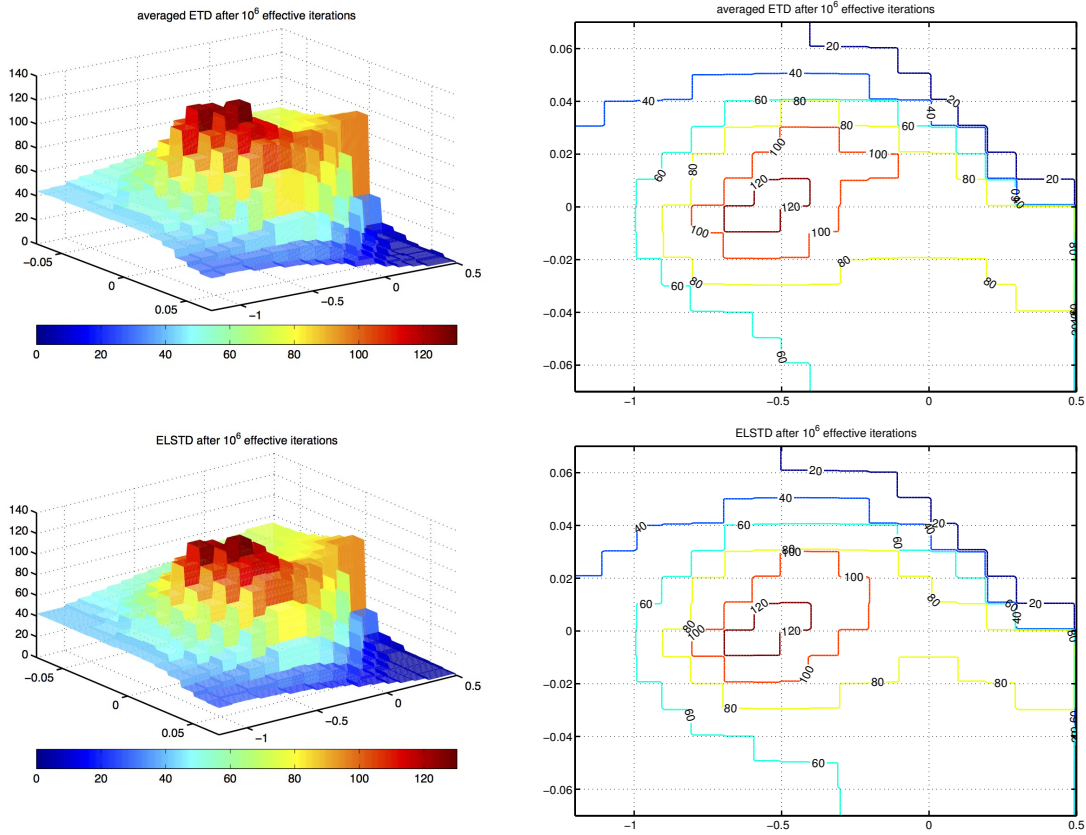


Figure 34: Piecewise constant approximation of $-v_\pi$ calculated by Variant I (top) and ELSTD (bottom) using tile-coding. Left: 3D view; right: contour map.

run, with averaging taken over the last 5×10^5 iterations as before to reduce transient effects. It can be seen that the results from Variant I and ELSTD (bottom row of Figure 34) are now much closer to each other than in Figure 33. Furthermore, both approximations can be compared with the approximate solution from the discretized model shown in Figure 31.

References

- [1] Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning*. MIT Press, Cambridge, MA.
- [2] Yu, H. (2012). Least squares temporal difference methods: An analysis under general conditions. *SIAM J. Control Optim.*, 50:3310–3343.
- [3] Yu, H. (2015). Weak convergence properties of constrained emphatic temporal-difference learning with constant and slowly diminishing stepsize. <http://arxiv.org/abs/1511.07471>.